

RECURRENCE AND TRANSIENCE FOR THE FROG MODEL ON TREES

CHRISTOPHER HOFFMAN, TOBIAS JOHNSON, AND MATTHEW JUNG

ABSTRACT. The frog model is a growing system of random walks where a particle is added whenever a new site is visited. A longstanding open question is how often the root is visited on the infinite d -ary tree. We prove the model undergoes a phase transition, finding it recurrent for $d = 2$ and transient for $d \geq 5$. Simulations suggest strong recurrence for $d = 2$, weak recurrence for $d = 3$, and transience for $d \geq 4$. Additionally, we prove a 0-1 law for all d -ary trees, and we exhibit a graph on which a 0-1 law does not hold.

1. INTRODUCTION

The frog model is a system of interacting random walks on a given rooted graph. Initially, the graph contains one particle at the root and some configuration of sleeping particles on its vertices; unless otherwise stated, we will assume an initial condition of one sleeping particle per vertex. The particle at the root starts out awake and performs a simple nearest-neighbor random walk in discrete time. Whenever a vertex with sleeping particles is first visited, all the particles at the site wake up and begin their own independent random walks, waking particles as they visit them. A formal definition of the frog model is in [AMP02a], and a nice survey of variations is in [Pop03]. Traditionally, particles are referred to as frogs, a practice we will continue here.

One of the most basic questions about the frog model on an infinite graph is whether it is recurrent or transient. Telcs and Wormald determined that the frog model was recurrent on \mathbb{Z}^d for any d , the first published result on the frog model [TW99]. On an infinite d -ary tree, this question is more difficult. It was first posed in [AMP02b]. It was asked again in [Pop03] and in [GS09], which pointed out that the answer was unknown even on a binary tree.

Our main result in this paper is that the frog model is recurrent on the binary tree but transient on the d -ary tree for $d \geq 5$, demonstrating a phase transition not found on \mathbb{Z}^d . A branching random walk martingale argument proves transience when $d \geq 6$. Pushing this result down to $d = 5$ is more complicated and requires computer assistance. Our proof of recurrence on the binary tree uses a bootstrapping argument in which we iteratively assume that the number of visits to the root is stochastically large and then prove it even larger; the argument seems novel to us.

Background on the frog model. It is common to use the frog model as a model for the spread of rumors or infections, thinking of awakened frogs as informed or infected agents. See [DG99] for an overview and [KPV04] for more tailored discussion. Another perspective on the frog model is as a conservative lattice gas model with the reaction $A + B \rightarrow 2A$. Here A represents an active particle and B an inert particle. Active particles disperse throughout the graph, igniting any inert particles they contact. Several papers taking this perspective study a process identical to the frog model except that particles move in continuous rather

than discrete time [RS04, CQR09, BR10]. This process and its variants have also seen much study by physicists; see the references in [CQR09] and [BR10]. Our results in this paper depend only on the paths of the frogs and not on the timing of their jumps, and so they apply equally well to this continuous-time process.

In the larger mathematical context, the frog model is part of a family of self-interacting random walks which have proven quite difficult to analyze. ([Pem07] provides a nice survey of this family.) In recent years progress on a few self-interacting random walks has generated considerable interest. One of these close relatives is activated random walk, which is touched on in [KS06] and studied in depth in [DRS10, RS12, ST14]. Activated random walk can be described as a frog model where frogs fall back asleep at some given rate. Another related process is excited random walk [BW03]. This walk has a bias the first time it visits a site but is unbiased each subsequent time that it returns. The frog model can be thought of as an “excited” branching process, which branches at a site v only the first time the process visits v .

Initial interest in the frog model was on the graph \mathbb{Z}^d . For any d , it was shown that the process is recurrent [TW99] and that the set of visited vertices grows linearly and when rescaled converges to a limiting shape [AMP02a]. A similar shape theorem was proven independently in [RS04] for the process with continuous-time particles. Both shape theorems rely on the subadditive ergodic theorem. A technical difficulty that arises is proving the expected time to wake a given frog is finite. Thus, measuring recurrence on a given graph is the first step in understanding the long time behavior of the model. Other results for the frog model on \mathbb{Z}^d include determining the decay of initially sleeping frogs necessary to make the model transient [Pop01]. Later this question was studied on the integers with drift [GS09]. A finer analysis in this setting recently appeared in [DP14] and [GNR15].

Our main interest in this paper is in recurrence and transience on \mathbb{T}_d , the infinite rooted d -ary tree. We denote the root by \emptyset . We also study aspects of the process on $\mathbb{T}_d^{\text{hom}}$, the infinite homogeneous degree $(d+1)$ -tree.

Some attention has already been given to a relative of our model on $\mathbb{T}_d^{\text{hom}}$ in which awake frogs die after independently taking a geometrically distributed number of jumps. In [AMP02b] and [LMP05], the authors prove a phase transition for survival. Depending on the parameter, there will either be frogs alive at all times with positive probability or the process will die out almost surely. We study the model in which frogs jump perpetually—a fundamentally different problem, since it switches the emphasis from the local to the global behavior of the model.

Statement and discussion of results. For a given rooted graph we call a realization of the model recurrent if the root is visited infinitely many times and transient if it is visited finitely often. Our main theorem covers the d -ary tree for all but two degrees:

Theorem 1.

- (i) *The frog model on \mathbb{T}_2 is almost surely recurrent.*
- (ii) *The frog model on \mathbb{T}_d for $d \geq 5$ is almost surely transient.*

We also make a conjecture on these two unsolved degrees based on fairly convincing evidence from simulations, presented in Section 5.

Conjecture 2. *The frog model on \mathbb{T}_d is recurrent a.s. for $d = 3$ and transient a.s. for $d = 4$.*

The simulations suggest the possibility of a three-phase transition as d increases. We call the model strongly recurrent if the probability the root is occupied is bounded away from

zero for all time. We call it weakly recurrent if it is recurrent with positive probability, but the probability that the root is occupied decays to zero.

Open Question 3. *Is the frog model strongly recurrent on \mathbb{T}_2 but weakly recurrent on \mathbb{T}_3 ?*

Such a transition would give information about the time to wake all children of the root. For instance, strong recurrence on \mathbb{T}_2 would imply that this time has finite expectation and an exponential tail. Should \mathbb{T}_3 exhibit a weak recurrence phase, then a tantalizing problem would be to estimate the decay of the occupation time at the root.

The recurrence of the frog model on the binary tree established in Theorem 1 (i) is the flagship result of this article. The proof goes by coupling the frog model with a process in which the root is visited less often. Let V be the number of visits to the root in this restricted model. The payoff is (2), a recursive distributional equation (RDE, see [AB05]) relating V and two independent copies of itself. We find that $V \sim \delta_\infty$ is the unique solution. Thus, the original model is recurrent.

The proof that $V \sim \delta_\infty$ is the unique solution of the RDE uses a seemingly novel bootstrapping argument. We assume that V dominates a Poisson and show that in fact, V dominates a Poisson with slightly larger mean. It follows by repeating this argument that V dominates any Poisson. One obstacle to making this work is that using the typical definition of stochastic dominance, we cannot establish a base case for the argument. To get around this, we instead use a weaker stochastic ordering defined in terms of generating functions, under which the argument holds even when starting with the trivial base case of V dominating the distribution $\text{Poi}(0)$. The situation is different for the frog model with initial conditions of $\text{Poi}(\mu)$ frogs per site. In this setting, we use the usual notion of stochastic dominance and a related bootstrapping argument to prove a phase transition from transience to recurrence on any d -ary tree as μ increases [HJJ15].

We believe that these ideas are more widely applicable. Aldous and Bandyopadhyay study RDEs in general in [AB05]. Another example of analyzing an RDE through an induced relation of generating functions can be found in [Liu98]. The RDE (2) in this paper is specific to our setting and much more complicated than the RDEs analyzed in either of these sources. Still, we think that our argument can be applied to other RDEs, including ones derived from similar interacting particle systems like activated random walk and the frog model with death.

For the transience part of Theorem 1, the idea is to dominate the frog model by a branching process. At the beginning of Section 3, we show in a few lines that a doubling branching random walk is transient on the 15-ary tree. A simple refinement in Proposition 18 improves this to $d \geq 6$. The case $d = 5$ uses a branching random walk with 27 particle types. This is significantly more complicated, and computing the transition probabilities requires computer assistance. Conceptually our approach could extend to a computer-assisted proof for transience when $d = 4$, but the demands of this theoretical program seem well beyond current processing power.

We present two other results besides Theorem 1. The first is a 0-1 law for transience and recurrence of the frog model on a d -ary tree that applies under more general initial conditions than one frog per site. For a given distribution ν on the nonnegative integers, we consider the frog model on a d -ary tree with the number of sleeping frogs on each vertex other than the root drawn independently from ν . The root initially contains one frog, which begins its life awake. Recall that when a site is visited for the first time, all sleeping frogs at that site are awoken. We refer to this as the frog model with i.i.d.- ν initial conditions. When $\nu = \delta_1$, this is the usual one-per-site frog model. This theorem complements the

0-1 law for recurrence proven in [GS09] in a frog model on \mathbb{Z} with drift. It also plays an important role in [HJJ15], where we use it to show that the probability of recurrence for the frog model on a d -ary tree with i.i.d.- $\text{Poi}(\mu)$ initial conditions jumps abruptly from 0 to 1 as μ increases.

Theorem 4. *The frog model on \mathbb{T}_d for any d and any i.i.d. initial conditions is recurrent with probability 0 or 1.*

Our final related result is that in contrast to the 0-1 law on \mathbb{T}_d , there is a graph on which the frog model is recurrent with probability strictly between 0 and 1.

Theorem 5. *Let G be the graph formed by merging the root of \mathbb{T}_6 and the origin of \mathbb{Z} into one vertex. The frog model on G has probability $0 < p < 1$ of being recurrent.*

We remark that [Pop01] exhibits a frog model without a 0-1 law on \mathbb{Z}^d . In their example the initial distribution of frogs decays in the distance from the origin.

A few of our proofs would be simplified by changing the setting from d -ary to homogeneous trees. However, we are interested in applying these results to finite trees, and the infinite d -ary tree is more natural to work with from that perspective. In any event, the techniques underlying our theorems can all be cleanly modified to prove similar statements about the homogeneous tree.

2. RECURRENCE FOR THE BINARY TREE

An outline of our proof is as follows. We start by defining a process that we call the self-similar frog model. A consequence of Proposition 7 is that the number of visits to the root in this model is smaller than in the original one. Thus it suffices to prove the self-similar frog model recurrent. To do this, we define the random variable V to be the number of returns to the root and set $f(x) = \mathbf{E}x^V$, the generating function of V . The self-similarity of our model established in Proposition 6 allows us to show in Proposition 9 that the generating function satisfies the relation $f = \mathcal{A}f$ for an explicit operator \mathcal{A} . In Lemma 10, we show that \mathcal{A} is monotone on a large class of functions. Combining this with $f \leq 1$ on $[0, 1]$, we get

$$f = \mathcal{A}^n f \leq \mathcal{A}^n 1,$$

and in Lemma 14, we prove that this converges to 0 as $n \rightarrow \infty$. This implies that $f \equiv 0$ and $V = \infty$ a.s.

This proof can be interpreted as an argument about stochastic orders. One can define a stochastic order by saying that if X and Y are nonnegative integer-valued random variables and $\mathbf{E}t^X \geq \mathbf{E}t^Y$ for $t \in (0, 1)$, then X is *smaller in the probability generating function order* than Y . This order and an equivalent one called the Laplace transform order are discussed in [SS07, Section 5.A]. From this perspective, each application of the operator \mathcal{A} shows that the distribution of V is slightly larger in this stochastic order.

2.1. The non-backtracking frog model. We will define the *non-backtracking frog model*, in which frogs move as random non-backtracking walks stopped at the root. More formally, we define the random non-backtracking walk $(X_n, n \geq 0)$ as a process taking values in \mathbb{T}_d , with $X_0 = x_0$. On its first step, the walk moves to a uniformly random neighbor of x_0 . At every subsequent step, it chooses uniformly from its neighbors other than the one from which it arrived. Let $T = \inf\{n \geq 1 : X_n = \emptyset\}$, taking this to be ∞ if the walk never visits \emptyset . Define the non-backtracking frog model by changing the frog's paths in the definition of the frog model from simple random walks to the stopped non-backtracking walks given




FIGURE 1. Conditional on v being visited, V and $V_{v'}$ are identically distributed in the self-similar model.

by $(X_{n \wedge T}, n \geq 0)$. Notice that the initial frog is never stopped in this model, and only one child of the root is ever visited. Call this child \emptyset' .

2.2. The self-similar frog model. We now make a further alteration to the non-backtracking frog model. When a frog x in this model is woken, if its first step away from the root takes it to a vertex that has already been visited by another frog, we stop x after this step. If multiple frogs visit a vertex for the first time in the same step, we stop all but one of them. We call the resulting model the *self-similar frog model*.

Let $V = V_{\emptyset'}$ be the number of visits to the root in the self-similar frog model. Note that only frogs initially sleeping in $\mathbb{T}_d(\emptyset')$, the subtree rooted at \emptyset' , have a chance of visiting the root. Suppose that vertex v is visited by a frog. Conditional on this, let v' be the child of v that the waking frog moves to next, and define $V_{v'}$ as the number of visits to v from frogs in $\mathbb{T}_d(v')$, the subtree rooted at v' . The following proposition explains how the self-similar frog model earns its name.

Proposition 6. *The distribution of $V_{v'}$ conditional on some frog visiting v and moving next to v' is equal to the (unconditioned) distribution of V .*

Proof. Let x be the frog that wakes vertex v and moves from there to v' . Besides x , all frogs that start outside of $\mathbb{T}_d(v')$ get stopped when they try to enter this subtree. Thus, from the time that v' is woken on, if we consider the model restricted to $\{v\} \cup \mathbb{T}_d(v')$, it looks identical to the original self-similar frog model (see Figure 1).

To turn this into a precise statement, consider the model from the time x visits v on. Ignore the frog initially at v . Freeze frogs when they visit v from $\mathbb{T}_d(v')$. Since no frogs ever enter $\mathbb{T}_d(v')$, the process depends only on the frogs initially in $\mathbb{T}_d(v')$ and the initial frog x , and relabeling vertices $\{v\} \cup \mathbb{T}_d(v')$ as $\{\emptyset\} \cup \mathbb{T}_d(\emptyset')$ in the obvious way produces a process identically distributed as the original self-similar frog model. Thus V and $V_{v'}$ are functionals of identically distributed processes. \square

2.3. Coupling the models. Suppose we wanted to couple a simple and a non-backtracking random walk starting from a vertex v on the homogeneous tree $\mathbb{T}_d^{\text{hom}}$. Almost surely, there is a unique geodesic from v to infinity that intersects the walk infinitely many times, obtained by trimming away the backtracking portions from the walk. By symmetry, this geodesic is a uniformly random non-backtracking walk on $\mathbb{T}_d^{\text{hom}}$, coupled so that its path is a subset of the simple random walk's path. If we were working on $\mathbb{T}_d^{\text{hom}}$ and not \mathbb{T}_d , we could couple the non-backtracking and usual frog models as desired by coupling each frog in this way. To address the asymmetry of \mathbb{T}_d at its root, our coupling of non-backtracking and normal frogs on \mathbb{T}_d will involve an intermediate coupling with a random walk on $\mathbb{T}_d^{\text{hom}}$.




FIGURE 2. V is the total number of visits to \emptyset in the self-similar process, V_v and V_u are the number of visits to \emptyset' from frogs originally in $\mathbb{T}_2(v)$ and $\mathbb{T}_2(u)$, respectively. In the self-similar model V, V_v , and $V_u \mid \{u \text{ is visited}\}$ are identically distributed

Proposition 7. *There is a coupling of the non-backtracking, the self-similar and the usual frog models so that the path of every non-backtracking (self-similar) frog path is a subset of the path of the corresponding frog in the usual model.*

Proof. First, we couple a non-backtracking walk to a simple random walk not on \mathbb{T}_d , but on $\mathbb{T}_d^{\text{hom}}$. Let $(Y_n, n \geq 0)$ be a simple random walk on $\mathbb{T}_d^{\text{hom}}$ starting at x_0 . This random walk diverges almost surely to infinity, and there is a unique geodesic from x_0 to the path's limit. Let $(X_n, n \geq 0)$ be the path of this geodesic. By the symmetry of $\mathbb{T}_d^{\text{hom}}$, the process (X_n) is a random non-backtracking walk from x_0 .

Next, we consider \mathbb{T}_d as a subset of $\mathbb{T}_d^{\text{hom}}$ and define a new random walk $(Z_n, n \geq 0)$ by modifying (Y_n) as follows. First, delete all excursions of (Y_n) away from \mathbb{T}_d . This might leave the walk sitting at the root for consecutive steps; if so, we replace all consecutive occurrences of the root by a single one. This results in either an infinite path on \mathbb{T}_d or a finite path on \mathbb{T}_d truncated at a visit to the root. In the second case, we extend the path by tacking on an independent simple random walk to its end. It follows from the independence of excursions in simple random walk that the resulting process (Z_n) is a simple random walk on \mathbb{T}_d .

Let T be the first time past 0 that (X_n) hits the root, or ∞ if it never does. By our construction, $\{X_0, \dots, X_T\} \subseteq \{Z_n, n \geq 0\}$. Thus we have coupled the stopped non-backtracking walks and simple random walks on \mathbb{T}_d . Coupling each frog in the non-backtracking frog model to the corresponding frog in the usual model gives the desired coupling between the non-backtracking and usual frog models. As the self-similar model is obtained by stopping frogs in the non-backtracking model, we obtain a coupling for it as well. \square

2.4. Generating function recursion. We now apply the self-similarity described in Proposition 6 to obtain a relation satisfied by the generating function for the number of visits to the root in the self-similar model.

Definition 8. *Let V be the number of visits to the root in the self-similar frog model on \mathbb{T}_2 . Define $f: [0, 1] \rightarrow [0, 1]$ by $f(x) = \mathbf{E}x^V$ with the convention that if $V = \infty$ a.s. then $f(1) = 0$.*

Proposition 9. *Define \mathcal{A} , an operator on functions on $[0, 1]$, by*

$$(1) \quad \mathcal{A}g(x) = \frac{x+2}{3}g\left(\frac{x+1}{2}\right)^2 + \frac{x+1}{3}g\left(\frac{x}{2}\right)\left(1 - g\left(\frac{x+1}{2}\right)\right).$$

The generating function f satisfies $f = \mathcal{A}f$.




FIGURE 3. Outcomes that would result in events A , B and C , respectively. The path of the frog at \varnothing' is red and the path of a frog from the subtree $\mathbb{T}(v)$ is blue.

Proof. If $\mathbf{P}[V = \infty] = 1$ then $f \equiv 0$. This is easily checked to be a fixed point of \mathcal{A} . So, for the remainder of the argument suppose that $\mathbf{P}[V = \infty] < 1$.

The frog at the root in the self-similar model follows a non-backtracking path and visits one of its children and then one of this child's children; call these vertices \varnothing' and v , respectively. We label the yet to be visited child u (see Figure 2). Define V_v and V_u to be the number of frogs which visit \varnothing' that were originally sleeping in $\mathbb{T}_2(v)$ and $\mathbb{T}_2(u)$, respectively.

Proposition 6 guarantees that, since v has been visited, the random variable V_v is distributed identically to V . Conditionally on u being visited, the random variable V_u is also distributed identically to V . In fact, because frogs outside of $\mathbb{T}_2(u)$ affect $\mathbb{T}_2(u)$ only by determining whether or not u is visited, we can express V_u as $V_u = \mathbf{1}\{u \text{ is visited}\}V'$, where V' is distributed as V and is independent of V_v . This yields a description of V in terms of a pair of independent copies of itself:

$$(2) \quad V = \underbrace{\mathbf{1}\{\text{frog at } \varnothing' \text{ visits } \varnothing\}}_{\text{term 1}} + \underbrace{\mathbf{1}\{u \text{ is visited}\}\text{Bin}(V', \frac{1}{2})}_{\text{term 2}} + \underbrace{\text{Bin}(V_v, \frac{1}{2})}_{\text{term 3}}.$$

Term 1 accounts for a possible visit to \varnothing by the frog started at \varnothing' . The conditional binomial distributions in terms 2 and 3 arise because each frog that visits \varnothing' from u or v has a $\frac{1}{2}$ chance of jumping back to \varnothing .

Despite the independence between V_v and V' , the three terms are dependent. For example, if term 1 is zero, then term 2 is more likely to be nonzero, since the frog at \varnothing' not visiting \varnothing makes it more likely to visit u . We unearth the pairwise independence of V_v and V' from (2) by conditioning on the following three disjoint events (see Figure 3):

- A. the frog starting at \varnothing' visits u ;
- B. the frog at \varnothing' does not visit u , and a frog returns to \varnothing' through v and visits u ;
- C. no frog ever visits u .

Event A occurs with probability $1/3$. Given that k frogs return to \varnothing' through v , the probability of C is $(2/3)2^{-k}$. Since the number of frogs returning to \varnothing' through v is distributed identically to V , the probability of C is $\frac{2}{3}\mathbf{E}(\frac{1}{2})^V$, which we call $2q/3$. The probability of event B is $2(1 - q)/3$. Note that under our assumption $\mathbf{P}[V = \infty] < 1$ it follows that $0 < q < 1$.

Conditional on event A , B , or C , the terms in (2) are independent. Indeed, conditioning on whether u is visited makes terms 2 and 3 independent, and conditioning further on whether the frog at \varnothing' visits u then makes term 1 independent of the other two. Now, we describe the distributions of each term in (2) conditional on events A , B , and C . For a

given random variable X , we use $\text{Bin}(X, p)$ to denote the random variable $\sum_{i=1}^X B_i$, where $\{B_i\}_{i \in \mathbb{N}}$ are distributed as $\text{Bernoulli}(p)$, independent of each other and of X .

- Conditional on A , term 1 is 0 and terms 2 and 3 are distributed as independent $\text{Bin}(V, 1/2)$.
- Conditional on B , term 1 is $\text{Bernoulli}(1/2)$, term 2 is $\text{Bin}(V, 1/2)$, and term 3 is $\text{Bin}(V, 1/2)$ conditional on being strictly less than V (since at least one frog will visit u and not move to \emptyset).
- Conditional on C , term 1 is $\text{Bernoulli}(1/2)$, term 2 is 0, and term 3 is $\text{Bin}(V, 1/2)$ conditional on being equal to V (since every frog counted by V_v will return to \emptyset).

To summarize, let X' and X be distributed as $\text{Bin}(V, 1/2)$. Let Y be distributed as $\text{Bin}(V, 1/2)$ conditional on $\text{Bin}(V, 1/2) < V$. Let Z be distributed as $\text{Bin}(V, 1/2)$ conditional on $\text{Bin}(V, 1/2) = V$. Let $I \sim \text{Bernoulli}(1/2)$. Take all of these to be independent. Conditioning on events A , B , and C , equation (2) yields

$$(3) \quad V \stackrel{d}{=} \begin{cases} X' + X & \text{with probability } 1/3, \\ I + X' + Y & \text{with probability } 2(1-q)/3, \\ I + Z & \text{with probability } 2q/3. \end{cases}$$

From this description of the distribution of V ,

$$(4) \quad \begin{aligned} \mathbf{E}x^V &= \frac{1}{3} \mathbf{E}x^{X'+X} + \frac{2(1-q)}{3} \mathbf{E}x^{I+X'+Y} + \frac{2q}{3} \mathbf{E}x^{I+Z} \\ &= \frac{1}{3} \mathbf{E}x^{X'} \mathbf{E}x^X + \frac{2(1-q)}{3} \mathbf{E}x^I \mathbf{E}x^{X'} \mathbf{E}x^Y + \frac{2q}{3} \mathbf{E}x^I \mathbf{E}x^Z. \end{aligned}$$

Recall that a $\text{Bernoulli}(p)$ random variable has generating function $px + 1 - p$ and that a random sum of i.i.d. random variables, $\sum_1^N X_i$, has generating function $g_N(g_{X_1}(x))$, where g_N and g_{X_1} are the generating functions of N and X_1 . From these facts,

$$\begin{aligned} \mathbf{E}x^I &= \frac{x+1}{2}, \\ \mathbf{E}x^{X'} &= \mathbf{E}x^X = f\left(\frac{x+1}{2}\right). \end{aligned}$$

The generating functions $\mathbf{E}x^Y$ and $\mathbf{E}x^Z$ are a bit more complicated. The random variable Y is distributed as X conditional on $X < V$. Using the basic formula for conditional probability,

$$\begin{aligned} \mathbf{P}[Y = k] &= \mathbf{P}[X = k \mid X < V] = \frac{\mathbf{P}[X = k \text{ and } X < V]}{\mathbf{P}[X < V]} \\ &= \frac{\mathbf{P}[X = k] - \mathbf{P}[X = V = k]}{1 - q} \\ &= \frac{\mathbf{P}[X = k] - 2^{-k} \mathbf{P}[V = k]}{1 - q}. \end{aligned}$$

Thus, the probability generating function of Y is

$$\begin{aligned}
 \mathbf{E}x^Y &= \frac{1}{1-q} \sum_{k=0}^{\infty} x^k (\mathbf{P}[X = k] - 2^{-k} \mathbf{P}[v = k]) \\
 (5) \quad &= \frac{1}{1-q} \mathbf{E} \left[x^X - \left(\frac{x}{2} \right)^V \right] \\
 &= \frac{1}{1-q} \left(f\left(\frac{x+1}{2}\right) - f\left(\frac{x}{2}\right) \right).
 \end{aligned}$$

In (5) we are making use of the general fact that $\sum (a_n - b_n) = \sum a_n - \sum b_n$ so long as each sum is finite. Similarly,

$$\mathbf{P}[Z = k] = \mathbf{P}[X = k \mid X = V] = \frac{2^{-k} \mathbf{P}[V = k]}{q},$$

and so

$$\mathbf{E}x^Z = \frac{1}{q} \sum_{k=0}^{\infty} x^k 2^{-k} \mathbf{P}[V = k] = \frac{1}{q} f\left(\frac{x}{2}\right).$$

Using all of these generating functions and (4)

$$\begin{aligned}
 f(x) &= \frac{1}{3} \mathbf{E}x^{X'} \mathbf{E}x^X + \frac{2(1-q)}{3} \mathbf{E}x^I \mathbf{E}x^{X'} \mathbf{E}x^Y + \frac{2q}{3} \mathbf{E}x^I \mathbf{E}x^Z \\
 &= \frac{1}{3} f\left(\frac{x+1}{2}\right)^2 + \frac{2(1-q)}{3} \left(\frac{x+1}{2} f\left(\frac{x+1}{2}\right) \frac{1}{1-q} \left(f\left(\frac{x+1}{2}\right) - f\left(\frac{x}{2}\right) \right) \right) \\
 &\quad + \frac{2q}{3} \left(\frac{x+1}{2q} f\left(\frac{x}{2}\right) \right) \\
 &= \frac{x+2}{3} f\left(\frac{x+1}{2}\right)^2 - \frac{x+1}{3} f\left(\frac{x+1}{2}\right) f\left(\frac{x}{2}\right) + \frac{x+1}{3} f\left(\frac{x}{2}\right) = \mathcal{A}f(x),
 \end{aligned}$$

which establishes our claim. \square

2.5. Proving recurrence. We have reduced the problem to understanding the properties of the operator \mathcal{A} defined in (1). In Lemma 10, we prove that \mathcal{A} is monotonic for functions belonging to the set $\mathcal{S} = \{g: [0, 1] \rightarrow [0, 1], \text{ nondecreasing}\}$. In Lemma 11, we show that \mathcal{A} maps \mathcal{S} into itself, so that we can apply Lemma 10 after applying \mathcal{A} iteratively. Finally, we show in Lemmas 12 and 14 that $\mathcal{A}^n 1 \rightarrow 0$. Starting at the conclusion of Proposition 9 (that the generating function f is a fixed point of \mathcal{A}), we will then apply these results to show that $f \equiv 0$, thus proving that the number of visits to the root in the self-similar frog model is a.s. infinite.

Lemma 10. *Let $g, h \in \mathcal{S}$. If $g \leq h$, then $\mathcal{A}g \leq \mathcal{A}h$.*

Proof. For $0 \leq t \leq 1$ define the interpolation between g and h by

$$i_t(x) = (1-t) \cdot g(x) + t \cdot h(x).$$

Since $\mathcal{A}i_0 = \mathcal{A}g$ and $\mathcal{A}i_1 = \mathcal{A}h$ it suffices to prove that $\frac{d}{dt} \mathcal{A}i_t(x) \geq 0$. Fix x and set $a = i_t\left(\frac{x+1}{2}\right)$ and $b = i_t\left(\frac{x}{2}\right)$ so that

$$\mathcal{A}i_t(x) = \frac{2+x}{3} a^2 + \frac{1+x}{3} b(1-a).$$

Define $s(a, b) = \mathcal{A}i_t(x)$. The chain rule implies

$$\frac{d}{dt}\mathcal{A}i_t(x) = \frac{\partial}{\partial a}s(a, b)\frac{da}{dt} + \frac{\partial}{\partial b}s(a, b)\frac{db}{dt}.$$

To prove $\frac{d}{dt}\mathcal{A}i_t \geq 0$ it suffices to prove each term in the above formula is nonnegative.

- The assumption that $g \leq h$ implies that $\frac{d}{dt}i_t(x) = h(x) - g(x) \geq 0$ for all t and x . In particular, this implies $\frac{da}{dt}, \frac{db}{dt} \geq 0$.
- First we compute the partials

$$\frac{\partial}{\partial a}s(a, b) = 2a\frac{2+x}{3} - b\frac{1+x}{3} \quad \text{and} \quad \frac{\partial}{\partial b}s(a, b) = (1-a)\frac{1+x}{3}.$$

As g and h are nondecreasing, i_t is also nondecreasing in x for any fixed t . Hence $b \leq a$. Along with the bound $a \leq 1$ this immediately implies both partials are positive. \square

Lemma 11. *If $g \in \mathcal{S}$, then $\mathcal{A}g \in \mathcal{S}$.*

Proof. All summands in (1) are nonnegative when $g(x) \leq 1$, which implies that $\mathcal{A}g \geq 0$. By the previous lemma, $\mathcal{A}g \leq \mathcal{A}1 \leq 1$. We can conclude then that $0 \leq \mathcal{A}g \leq 1$. To see that $\mathcal{A}g$ is nondecreasing, suppose that $x \leq y$, and let $a = g(\frac{y+1}{2}) - g(\frac{x+1}{2})$. Then we have

$$\begin{aligned} \mathcal{A}g(y) &\geq \frac{x+2}{3}g\left(\frac{x+1}{2}\right)g\left(\frac{y+1}{2}\right) + \frac{x+1}{3}g\left(\frac{x}{2}\right)\left(1 - g\left(\frac{y+1}{2}\right)\right) \\ &= \mathcal{A}g(x) + \left(\frac{x+2}{3}g\left(\frac{x+1}{2}\right) - \frac{x+1}{3}g\left(\frac{x}{2}\right)\right)a \geq \mathcal{A}g(x). \end{aligned} \quad \square$$

We now analyze the behavior of \mathcal{A} on the family of generating functions for Poisson random variables. Recall that the generating function of $\text{Poi}(a)$ is $e^{a(x-1)}$.

Lemma 12. *Define $g_a(x) = e^{a(x-1)}$ for all $a \geq 0$. For all $x \in [0, 1]$,*

$$\mathcal{A}g_a(x) \leq g_{a+c_a}(x),$$

where

$$(6) \quad c_a = \begin{cases} \frac{1}{3}e^{-2} & 0 \leq a \leq 4, \\ \frac{1}{3}e^{-a/2} & a \geq 4. \end{cases}$$

Proof. Applying the operator \mathcal{A} , we have

$$(7) \quad \begin{aligned} \mathcal{A}g_a(x) &= \frac{x+2}{3}e^{a(x-1)} + \frac{x+1}{3}e^{ax/2-a}\left(1 - e^{a(x-1)/2}\right) \\ &= g_a(x)r_{a/2}(x), \end{aligned}$$

where

$$r_b(x) = \frac{2+x}{3} + \frac{1+x}{3}(e^{-bx} - e^{-b}).$$

Note that $g_a(x)g_b(x) = g_{a+b}(x)$. It thus suffices to establish

Claim. *For $x \in [0, 1]$, we have $r_b(x) \leq g_{c_{2b}}(x)$.*

Proof of claim. We drop subscripts and let $r(x) = r_b(x)$ and $c = c_{2b}$. Calculus and a little algebra show that

$$r'(x) = \frac{1}{3} (1 - e^{-b} + e^{-bx}(-bx - b + 1))$$

and

$$r''(x) = \frac{1}{3} e^{-bx} (b^2(x+1) - 2b).$$

We break the proof up into cases.

- If $b \leq 1$ then $r(x)$ is concave down on $[0, 1]$ and the graph of $r(x)$ lies below its tangent line at $x = 1$. Thus

$$\begin{aligned} r(x) &\leq 1 + r'(1)(x-1) = 1 + \frac{1}{3} [1 - 2be^{-b}](x-1) \\ &\leq \exp \left[\frac{1}{3} (1 - 2be^{-b})(x-1) \right]. \end{aligned}$$

It is easily verified that $\frac{1}{3}(1 - 2be^{-b}) \geq \frac{1}{3}(1 - 2e^{-1}) \geq e^{-2}/3$ for $b \leq 1$ and hence that $r(x) \leq g_c(x)$.

- If $b \geq 2$ then $r(x)$ is concave up on $[0, 1]$ and the graph of $r(x)$ lies below the secant line between $(0, r(0))$ and $(1, r(1))$. Thus as $r(1) = 1$ we have

$$\begin{aligned} r(x) &\leq 1 + (1 - r(0))(x-1) = 1 + \frac{1}{3} e^{-b}(x-1) \\ &\leq \exp \left[\frac{1}{3} e^{-b}(x-1) \right] = g_c(x). \end{aligned}$$

- If $1 < b < 2$ then there is a unique inflection point at $I = \frac{2}{b} - 1$ where r switches from concave down to concave up. Since r is concave up on $[I, 1]$, the graph of r lies below the line connecting $(1, 1)$ to $(I, r(I))$. Since r is concave down on $[0, I]$, to the left of I the graph of r lies below its tangent line at $(I, r(I))$. Thus the line segment from $(I, r(I))$ to $(0, r(I) - r'(I)I)$ lies above r , as in Figure 4. Therefore r lies below the line between $(1, 1)$ and $(0, r(I) - r'(I)I)$, and

$$(8) \quad r(x) \leq 1 + (1 - r(I) + Ir'(I))(x-1).$$

Next, we evaluate

$$(9) \quad 1 - r(I) + Ir'(I) = 1 - \frac{1}{3} \left(2 + \left(\frac{4}{b} - 1 \right) e^{b-2} - e^{-b} \right)$$

and try to bound this expression from below for $b \in (1, 2)$. We proceed as calculus students, looking for critical points in this interval. The derivative with respect to b is

$$-\frac{1}{3} \left(\left(\frac{4}{b} - 1 - \frac{4}{b^2} \right) e^{b-2} + e^{-b} \right),$$

and a bit of algebra shows that the zeros of this expression are the solutions to

$$e^{2(b-1)} \left(\frac{2-b}{b} \right)^2 = 1.$$

Taking logarithms, we are interested in solutions to

$$b - 1 + \log(2 - b) - \log b = 0.$$




FIGURE 4. Above the graph of $y = r(x)$ sits the secant line from $(I, r(I))$ to $(1, 1)$ and the tangent line to $r(x)$ at $x = I$, both depicted in blue. Above them in red is the line $y = 1 + (1 - r(I) + Ir'(I))(x - 1)$.

on $(1, 2)$. On this interval we can replace the logarithms with their power series expansions around 1 to rewrite the left-hand side as

$$b - 1 + 2 \left(\frac{(b-1)^2}{2} + \frac{(b-1)^4}{4} + \frac{(b-1)^6}{6} + \dots \right),$$

which is strictly positive for $b \in (1, 2)$. Thus (9) has no critical values on $(1, 2)$, and its minimum on $[1, 2]$ is $e^{-2}/3$, occurring at $b = 2$. Applying this to (8), we have shown that

$$r(x) \leq 1 + \frac{1}{3}e^{-2}(x-1) \leq \exp \left[\frac{1}{3}e^{-2}(x-1) \right] = g_c(x).$$

This concludes the proof of both the claim and the lemma. \square

Remark 13. Though the preceding lemma was an exercise in calculus, it has a probabilistic interpretation. If we think of \mathcal{A} as acting directly on distributions instead of on their generating functions, this lemma shows that the result of applying \mathcal{A} to $\text{Poi}(a)$ is larger than $\text{Poi}(a + c_a)$ in the probability generating function stochastic order described at the beginning of Section 2. The reason that $\mathcal{A}g_a$ simplifies so nicely in (7) is the Poisson thinning property, and the fact that $g_a(x)g_b(x) = g_{a+b}(x)$ is just the statement that the sum of independent Poissons is Poisson. There is a temptation to interpret $\mathcal{A}g_a(x) = g_a(x)r_{a/2}(x)$ as saying that the distribution resulting from applying \mathcal{A} to $\text{Poi}(a)$ is a convolution of $\text{Poi}(a)$ and another distribution, but $r_{a/2}(x)$ is not monotone in x and hence not the generating function of a probability distribution.

Lemma 14. For $x \in [0, 1)$,

$$\lim_{n \rightarrow \infty} \mathcal{A}^n g_0(x) = 0.$$

Proof. Define the sequence a_n by $a_0 = 0$ and $a_{n+1} = a_n + c_{a_n}$. By Lemmas 10, 11, and 12,

$$\mathcal{A}^n g_0(x) \leq g_{a_n}(x) = e^{a_n(x-1)}.$$

We need to show that $a_n \rightarrow \infty$ as $n \rightarrow \infty$. Suppose this does not hold. Since the sequence is increasing, $a_n \rightarrow a$ for some constant a . Looking back at (6), this implies that c_{a_n} converges to a strictly positive limit. We can then choose n sufficiently large that $a_n + c_{a_n} > a$, a contradiction. \square

Proof of Theorem 1 (i). Let f be the generating function $f(x) = \mathbf{E}x^V$ with V the number of visits to the root in the self-similar model frog model on the binary tree. By Proposition 9 we know that f satisfies the recursion relation $\mathcal{A}f = f$. Since f is a probability generating function, it satisfies $f(x) \leq 1 = g_0(x)$ for $x \in [0, 1]$. Proposition 9 and Lemmas 10 and 11 imply $f(x) \leq \mathcal{A}^n g_0(x)$ for all n . By Lemma 14, f is identically zero on $[0, 1)$. Thus the probability of any finite number of returns to the root is 0. This implies there are a.s. infinitely many returns to the root in the self-similar model. By the coupling in Proposition 7 each return in the self-similar model corresponds to a distinct return in the frog model. So, the frog model on the binary tree is a.s. recurrent. \square

3. TRANSIENCE FOR $d \geq 5$

The non-backtracking model was useful in the previous section because it was dominated by the usual frog model but was still recurrent. To prove transience we instead seek processes that dominate the frog model and can be proven transient. For example, consider a branching random walk on \mathbb{T}_d whose particles split in two at every step. Using a union bound and asymptotics for the number of Dyck paths, one can show that the probability that any of the 2^{2n} particles at time $2n$ are at the root is $O(n^{-3/2})$ when $d \geq 15$, and hence the branching random walk visits the root finitely many times. As this walk can be naturally coupled to the frog model so that every awake frog has a corresponding particle, this proves that the frog model is transient for $d \geq 15$.

In this section, we will present a series of refinements to this argument to ultimately prove Theorem 1 (ii). In Proposition 18, we use a branching random walk on the integers and martingale techniques to prove transience for $d \geq 6$. We use this argument as a base for our proofs of Proposition 19, transience on the deterministic tree which alternates between five and six children, and Theorem 1 (ii), transience for $d \geq 5$. Both proofs use a multitype branching random walk. We included Proposition 19 because its calculations can easily be done by hand. In Theorem 1 (ii), on the other hand, we use a branching random walk with 27 types. The necessary calculations are intractable by hand, but they take only a few seconds on a computer. To get started we first address some difficulties that arise from reflection at the root. In doing so we also prove the 0-1 law described in Theorem 4.

3.1. Couplings and 0-1 law. We will need to consider the frog model on several modifications of a rooted tree. We can handle these special cases all at once by working in a more general setting. Let Λ be any infinite rooted graph and H any graph. Enumerate finitely or countably many copies of Λ by $\Lambda(i)$, and form a graph G by adding an edge from the root of each $\Lambda(i)$ into H . Our next lemma shows that regardless of the number of sleeping frogs placed on H , a frog model is less transient on G than on Λ . Our motivation is the case when $\Lambda = \mathbb{T}_d$, as in Corollaries 16 and 17.




FIGURE 5. The paths of frog x in Λ and frog x' in G . Frog x follows the blue steps of x and ignores the red steps.

Lemma 15. *We consider two frog models. The first is on Λ with i.i.d.- ν initial conditions, for any measure ν on the nonnegative integers. The second is on G with the following initial conditions: one initially active frog at the root of $\Lambda(1)$; i.i.d.- ν sleeping frogs at all other vertices of $\bigcup_i \Lambda(i)$; and any configuration of sleeping frogs in H . Assume H is such that a random walk on G a.s. escapes H .*

Let V_G be the number of times the root of any $\Lambda(i)$ is visited in the frog model on G , not counting steps from H to a root. Let V_Λ be the the number of times the root is visited in the frog model on Λ . Then the two frog models can be coupled so that $V_G \geq V_\Lambda$.

Proof. Have the frog x , awake at the root $\emptyset \in \Lambda$, mime the frog x' that starts at the root of $\Lambda(1)$. As depicted in Figure 5, whenever x' enters H the frog x pauses at \emptyset ; when x' re-enters any $\Lambda(i)$, the frog x begins following x' again.

When x visits a vertex that has yet to be visited, so will x' . Couple the number of sleeping frogs at the vertices occupied by x and x' , and couple the newly awoken frogs to each other as with x and x' . In this way, x and all descendants on Λ perform simple random walks coupled to a frog on some $\Lambda(i)$. Thus, every visit to the root in Λ corresponds to a visit to level 0 in G , showing that $V_G \geq V_\Lambda$ under this coupling. \square

We give two corollaries. The first will help us prove our transience results, and the second will help us prove a 0-1 law for transience and recurrence.

Corollary 16. *Consider the frog model on the $(d+1)$ -homogeneous tree $\mathbb{T}_d^{\text{hom}}$ starting with a single active frog at the root, and with no sleeping frog at direct ancestors of the root. If level 0 is almost surely visited finitely many times in this model, then the frog model on \mathbb{T}_d is almost surely transient.*

Proof. Let $G = \mathbb{T}_d^{\text{hom}}$, thinking of it as countably many copies of \mathbb{T}_d each joined at its root to a leaf of the infinite graph consisting of all the negative levels of G . The statement then follows immediately from Lemma 15. \square

Corollary 17. *Run the frog model on \mathbb{T}_d , starting with an active frog not at the root but at level k . Assume that there are no sleeping frogs at levels $0, \dots, k-1$ and i.i.d.- ν sleeping frogs at level k and beyond, with the exception of the location of the initial frog. The probability that the root is visited infinitely often in this model is at least the probability that the root is visited infinitely often in the usual frog model on \mathbb{T}_d with i.i.d.- ν initial conditions.*

Proof. To set up our alternate frog model, let $G = \mathbb{T}_d$, thinking of it as d^k copies of \mathbb{T}_d joined by a graph consisting of levels 0 to $k-1$ of the original graph. Let p be the probability that the root is visited infinitely often in the usual frog model. Let Y be the number of visits from level $k+1$ to k in the alternate model. It follows immediately from Lemma 15 that $\mathbf{P}[Y = \infty] \geq p$.

Let X be the number of visits to the root. We would like to show that

$$(10) \quad \mathbf{P}[Y = \infty, X < \infty] = 0,$$

thus proving that $\mathbf{P}[X = \infty] \geq p$. Call it a *dash* if a frog moves from level $k+1$ to a vertex v at level k , walks directly to the root, and then walks directly back to v . Let X' be the total number of dashes that occur. Conditional on a frog stepping from level $k+1$ to k , it makes a dash independently of all other frogs, since the model has no sleeping frogs at levels 0 to $k-1$. Whether or not it makes a dash is also independent of its own future number of visits from level $k+1$ to k and of dashes. Thus, at every visit from level $k+1$ to k , there is an independent $1/d(d+1)^{2k-1}$ chance of a dash, showing that

$$\mathbf{P}[Y = \infty, X' < \infty] = 0.$$

Since $X' < X$, this shows (10) and completes the proof. \square

We are now ready to prove the 0-1 law.

Proof of Theorem 4. Suppose the probability that the root is visited infinitely often in the frog model on \mathbb{T}_d with i.i.d.- ν initial conditions is $p > 0$. We wish to show that $p = 1$. The idea of the proof is to turn this statement into a more finite event, and then show that there are infinitely many independent opportunities for this event to occur. To this end, fix a constant N . We will show that at least N frogs visit the root with probability 1.

Claim. *For any k and N , there is a constant $K = K(k, N)$ such that the following statement holds. Consider the frog model on \mathbb{T}_d starting with a frog at level k , with i.i.d.- ν sleeping frogs at levels $k, k+1, \dots, K-1$ with the exception of the vertex of the initial frog, and with no sleeping frogs outside of this range. With probability at least $p/2$, this process makes at least N visits to the root.*

Proof. Consider the frog process with no sleeping frogs below level k , as in Corollary 17. Let E_K be the event that there are at least N visits to the root by frogs that are woken without the help of any frogs at level K or beyond. As $K \rightarrow \infty$, the event E_K converges upward to the event that there are at least N visits to the root by any active frog, which occurs with at least probability p by Corollary 17. Thus, for sufficiently large K , we have $\mathbf{P}[E_K] \geq p/2$. \square

Now, we can find infinitely many independent events with probability $p/2$, each implying N visits to the root. Let $k_0, k'_0 = 0$, and inductively choose k_i, k'_i as follows. Let $k'_i = K(k_{i-1}, N)$ from the claim. Let k_i be the level of the first frog that wakes up at level k'_i or beyond (assuming that ν is not a point mass at 0, there will be such a frog). Now, imagine a frog process starting with this frog, with no sleeping frogs below level k_i or at level k_{i+1}

or beyond. These processes can all be embedded into the original frog process on \mathbb{T}_d , and each one independently has a $p/2$ chance of visiting the root at least N times, by the claim. Thus, the root is visited at least N times almost surely, for arbitrary N . \square

3.2. Proving transience. Consider the branching random walk where each particle gives birth at each step either to one child to its left or to two children to its right. Formally, we define this as a sequence of point processes. Start with ξ_0 as a single particle at 0. With probability $1/(d+1)$, the point process ξ_1 consists of a single particle at -1 ; with probability $d/(d+1)$, it consists of two particles at 1. After this, each particle in ξ_n produces children in ξ_{n+1} in the same way relative to its position, independently of all other particles. We will use this branching random walk to prove the frog model transient for $d \geq 6$ and closely related processes to extend this down to $d = 5$.

Proposition 18. *For $d \geq 6$, the frog model on \mathbb{T}_d is almost surely transient.*

Proof. Consider the frog model on $\mathbb{T}_d^{\text{hom}}$, starting with no sleeping frogs at direct ancestors of the root, as in Corollary 16. When a frog jumps backward in this process, it never spawns a new frog, and when it moves forward, it sometimes does. Thus, the projection of this frog model onto the integers can be coupled with $(\xi_n, n \geq 0)$ so that every frog has a corresponding particle. By Corollary 16, proving that ξ_n visits 0 finitely many times a.s. proves that the frog model on \mathbb{T}_d is transient a.s.

To determine the behavior of ξ_n , we define a weight function w on point process configurations. We refer to the position of a particle i in a point process configuration by $P(i)$ and define

$$(11) \quad w(\xi) = \sum_{i \in \xi} e^{-\theta P(i)},$$

with θ to be chosen later. Letting $\mu = \mathbf{E}w(\xi_1)$ we have

$$\mathbf{E}[w(\xi_{n+1}) \mid \xi_n] = \mu w(\xi_n),$$

and so the sequence $w(\xi_n)/\mu^n$ is a martingale. As it is positive, it converges almost surely. When $\mu < 1$ this means $w(\xi_n) \rightarrow 0$. If a particle in ξ_n occupies the origin then $w(\xi_n) \geq 1$, and so infinitely many visits to the origin prevents $w(\xi_n)$ from converging. Hence, $\mu < 1$ implies the a.s. transience of ξ_n . (In fact, this holds when $\mu = 1$ as well, though we will not need this.) It then suffices to show that there exists θ making $\mu < 1$. We compute

$$\mu = \frac{1}{d+1}e^\theta + 2\frac{d}{d+1}e^{-\theta}.$$

This is minimized by setting $\theta = \log(2d)/2$, which makes $\mu = 2\sqrt{2d}/(d+1)$. A bit of algebra shows that $\mu < 1$ when $d > 3 + 2\sqrt{2} \approx 5.83$. \square

By using a multitype branching process, we can extend this proof to show transience for \mathbb{T}_5 . Before we do so, we will show how it works in a setting where humans can do the math without much assistance.

Proposition 19. *Let $\mathbb{T}_{5,6}$ be the tree whose levels alternate between vertices with 5 children and vertices with 6 children, starting with the root having either 5 or 6 children. The frog model on this tree is transient a.s.*

Proof. Let $\mathbb{T}_{5,6}^{\text{hom}}$ be the five-six children alternating homogeneous tree which contains $\mathbb{T}_{5,6}$ and place a sleeping frog at each vertex except for direct ancestors of the root of $\mathbb{T}_{5,6}$. Lemma 15 implies that it suffices to prove transience of this frog model on $\mathbb{T}_{5,6}^{\text{hom}}$.

First note that a frog at a vertex with five children has different probabilities of moving forwards or backwards than a frog at a vertex with six children. By design the tree deterministically alternates, so a frog also alternates between each state.

When a frog moves backwards there is chance it immediately jumps forward to the same vertex, which will never spawn a new frog. Similarly, when two frogs occupy the same site there is a chance both jump forward to the same vertex, spawning at most one frog, not two. The idea is to introduce additional particle types that act like frogs in these more advantageous states.

Consider a multitype branching random walk on \mathbb{Z} with six particle types, F_5 , D_5 , B_5 , F_6 , D_6 , and B_6 . The subscript accounts for whether a frog is at a vertex with 5 or 6 children. B particles represent frogs that have just stepped backward. D particles represent two frogs at once, the waker and wakee at a vertex where a frog has just woken up. Last, F particles represent single frogs with sleeping frogs present at all children. A visual depiction of these particle types is provided in Figure 6, and the distribution of children for each particle type is defined in Figure 7.

Let ζ_n be the branching random walk in which particles reproduce independently with the given child distributions. These distributions are chosen to match how the projections of frogs on the integers behave. Ignoring for a moment whether a frog is at a site with five or six children, when a frog jumps back it becomes of type B and when a new frog wakes it and its waker consolidate into a type D particle. Any extra frogs become a type F particle. These particles then reproduce independently on a “fresh” tree configured so that the particles always generate at least as many frogs as the projection of the actual frog model. For this reason we can couple the integer projection of the frog model on $\mathbb{T}_{5,6}^{\text{hom}}$ with ζ_n so that the particles representing awake frogs are a subset of ζ_n . It therefore suffices to prove that ζ_n is transient.

To analyze ζ_n , we use a generalization of the martingale from Proposition 18 to the multitype setting, introduced in [Big76]. Let $\zeta_n = \sum_i \zeta_n^i$, where i ranges over the six particle types and ζ_n^i denotes the restriction of ζ_n to particles of type i . Recalling the weight function w given by (11), we define a matrix $\Phi(\theta)$ by

$$\Phi_{ij}(\theta) = \mathbf{E}_i \left[w(\zeta_1^j) \right].$$

Here, we use \mathbf{E}_i to denote expectation when ζ_0 is a single particle at the origin of type i . Let w_n denote a row vector whose i th entry is $w(\zeta_n^i)$. Then

$$(12) \quad \mathbf{E}[w_{n+1} \mid \zeta_n] = w_n \Phi(\theta).$$

Thus, for any eigenvalue λ and associated right eigenvector v of $\Phi(\theta)$,

$$\mathbf{E}[w_{n+1} v \mid \zeta_n] = w_n \Phi(\theta) v = \lambda w_n v,$$

and so $w_n v / \lambda^n$ is a martingale.

Since $\Phi(\theta)$ is a nonnegative irreducible matrix, there is a positive eigenvalue $\phi(\theta)$ equal to the spectral radius of $\Phi(\theta)$ by the Perron–Frobenius theorem. The eigenvector $v(\theta)$ associated with $\phi(\theta)$ has strictly positive entries. We then have a positive martingale $w_n v(\theta) / \phi(\theta)^n$. If $\phi(\theta) < 1$, then it follows as in Proposition 18 that the branching random walk visits 0 finitely often, thus proving that the frog model is almost surely transient.




FIGURE 6. A depiction of the six particle types from the proof of Proposition 19. Each asterisk is a frog represented by the particle. The symbol \square signifies a vertex with a sleeping frog, and the symbol \square represents a vertex with no sleeping frog.




FIGURE 7. The distribution of children for each particle type in the proof of Proposition 19.

All that remains is to find some value of θ such that $\phi(\theta) < 1$. Ordering the rows and columns $F_5, D_5, B_5, F_6, D_6, B_6$ and reading off $\mathbf{E}_i[w(\zeta_i^j)]$ from Figure 7,

$$\Phi(\theta) = \begin{bmatrix} 0 & 0 & 0 & 0 & \frac{5}{6}e^{-\theta} & \frac{1}{6}e^{\theta} \\ 0 & 0 & 0 & \frac{5}{36}e^{-\theta} & \frac{1}{36}e^{\theta} + \frac{55}{36}e^{-\theta} & \frac{5}{18}e^{\theta} \\ 0 & 0 & 0 & \frac{1}{6}e^{-\theta} & \frac{2}{3}e^{-\theta} & \frac{1}{6}e^{\theta} \\ 0 & \frac{6}{7}e^{-\theta} & \frac{1}{7}e^{\theta} & 0 & 0 & 0 \\ \frac{6}{49}e^{-\theta} & \frac{1}{49}e^{\theta} + \frac{78}{49}e^{-\theta} & \frac{12}{49}e^{\theta} & 0 & 0 & 0 \\ \frac{1}{7}e^{-\theta} & \frac{5}{7}e^{-\theta} & \frac{1}{7}e^{\theta} & 0 & 0 & 0 \end{bmatrix}.$$

Computing the eigenvalues of this matrix numerically, one can confirm that there exists θ with $\phi(\theta) < 1$; for example, $\phi(\log 3) \approx 0.9937$. To be completely certain that this is not an artifact of rounding, we will justify that $\phi(\log 3) < 1$ without using floating-point arithmetic. Observe that $\Phi(\log 3)$ has rational entries. Using the computer algebra system SAGE, we calculated $(\Phi(\log 3))^{66}$ using exact arithmetic, and we found that its largest row sum was less than 1. (The only significance of the 66th power is that it is the lowest one for which this is true.) This implies that all eigenvalues of $(\Phi(\log 3))^{66}$ are less than 1, which implies that all eigenvalues of $\Phi(\log 3)$ are less than 1 as well. The source code accompanying this paper includes this matrix and has instructions so that readers can easily check these claims. \square

Remark 20. We chose to include this proof to illustrate the technique we use to prove Theorem 1 (ii). Furthermore, this provides an example of proving the frog model is transient on an interpolation between different degree trees. This is relevant because the sharpest proof of Conjecture 2 would find exactly where the phase transition occurs between recurrence and transience on \mathbb{T}_d , perhaps between $d = 3$ and $d = 4$. Last, a natural generalization is a frog model on Galton-Watson trees. Our argument depends on the deterministic structure of $\mathbb{T}_{5,6}$ and we do not see an obvious way to generalize it.

Having proven transience for the frog model on \mathbb{T}_d with $d \geq 6$ and on $\mathbb{T}_{5,6}$, we present our final refinement to prove the \mathbb{T}_5 case. The proof is essentially the same as the previous one, but with more particle types and a more difficult calculation.

Proof of Theorem 1 (ii). We define a particle type $P(a, b, c)$, for $a \geq 1$ and $b, c \geq 0$. A particle of type $P(a, b, c)$ represents a frogs on one vertex. There are no sleeping frogs on at least b of the vertex's children and on at least c of the vertex's siblings. In this scheme, the F types from the previous proof would translate to $P(1, 0, 0)$, the D types would translate to $P(2, 0, 0)$, and the B types would translate to $P(1, 1, 0)$.

We use 27 of these particles, $P(a, b, c)$ with $1 \leq a \leq 3$ and $0 \leq b, c \leq 2$. For particle type $P(a, b, c)$, consider the frog model on the homogeneous tree, starting with a frogs at the root. As usual, remove the sleeping frogs from direct ancestors of the root. Also remove the sleeping frogs from b children of the root and from c siblings. From each of these 27 initial states, we compute all possible states to which the frog model could transition in two steps, along with the exact probabilities of doing so. We then represent each of these final states as a collection of particles of the 27 types, at levels $-2, 0$, and 2 on the tree. In this way, we determine child distributions for each particle type, as in Figure 7. There is a slight ambiguity in how to do this, as a state of frogs can be represented in more than one way by these particle types. For example, four frogs on one vertex with one sibling vertex with no sleeping frog could be represented as two particles of type $P(2, 0, 1)$, or as one of type $P(3, 0, 1)$ and one of type $P(1, 0, 1)$. We always chose particles greedily, opting for as many

3-frog particles as possible. Whatever choice we make here, our branching random walk will still dominate the frog model, since when we assign new particles we “reset” the tree below them so the particles wake at least as many frogs as their counterpart in the frog model.

As in Proposition 19, it suffices to compute the matrix $\Phi(\theta)$ and show that for some choice of θ , its highest eigenvalue is less than one. Our attached source code computes $\Phi(\theta)$ exactly. We include additional documentation there explaining how we performed this calculation and describing the steps we took to make sure it was trustworthy. To avoid rounding issues, we proceeded as with Proposition 19. We exactly computed $(\Phi(\log 3))^{1024}$ by successively squaring the matrix ten times, and we then checked that all of its row sums were less than 1. (There is no significance to the value $\log 3$; it just happens to work.) Thus, all eigenvalues of $(\Phi(\log 3))^{1024}$ are less than 1, implying that all eigenvalues of $\Phi(\log 3)$ are less than 1 as well. \square

4. A FROG MODEL WITHOUT A 0-1 LAW

We obtain a graph on which the frog model does not satisfy a 0-1 law by combining the transient graph \mathbb{T}_6 with the recurrent graph \mathbb{Z} and proving that there is a positive probability that the frogs in each do not interact much.

To this end, we first prove two lemmas. The first shows there is a positive probability that the rightmost frog on the \mathbb{Z} part of the combined graph escapes to ∞ while avoiding 0. This is necessary to rule out the possibility that too many frogs from \mathbb{Z} get lost forever in \mathbb{T}_6 . The second lemma proves there is a positive probability a frog model on \mathbb{T}_6 never returns to the origin, thus establishing a chance that the frog model on G gets lost in the transience of \mathbb{T}_6 .

Lemma 21. *Consider the frog model on G , the graph formed by merging the root of \mathbb{T}_6 and the origin of \mathbb{Z} . With positive probability, the frogs starting at $1, 2, \dots$ in \mathbb{Z} all wake up.*

Proof. Let

$$\delta_n = \frac{1}{8} \prod_{k=1}^{n-1} \left(1 - \frac{1}{(k+1)^2}\right),$$

taking $\delta_1 = 1/8$. We will show by induction that the frogs at $1, \dots, n$ wake up with probability at least δ_n . When $n = 1$, this holds because the initial frog moves right on its first step with probability $1/8$. Now, assume the statement for n . Condition on the frogs at $1, \dots, n$ being woken. From the time when the frog at n is woken on, the two frogs there are independent random walkers, and at least one of them reaches $n+1$ before 0 with probability $1 - 1/(n+1)^2$ by a standard martingale argument. Thus frog $n+1$ is woken with probability at least $\delta_n(1 - 1/(n+1)^2) = \delta_{n+1}$, completing the induction.

Taking a limit of increasing events, the probability of the frogs at $1, 2, \dots$ all waking is at least $\lim_{n \rightarrow \infty} \delta_n > 0$. \square

Lemma 22. *Let p' be the probability that the root is never visited past the initial frog's first move in the frog model on \mathbb{T}_6 . It holds that $p' > 0$.*

Proof. As in Proposition 18, consider the frog model on $\mathbb{T}_6^{\text{hom}}$, starting with no sleeping frogs at direct ancestors of the root. By Lemma 15, following the reasoning of Corollary 16, there is a coupling so that the number of visits to level 0 in the frog model on $\mathbb{T}_6^{\text{hom}}$ is at least the number of visits to the root in the frog model on \mathbb{T}_6 .

Now, recall from Proposition 18 the point process ξ , a branching random walk on \mathbb{Z} in which particles split whenever they move in the positive direction. This process dominates

the projection of the frog model on $\mathbb{T}_6^{\text{hom}}$ onto the integers. Putting this all together, it suffices to show that with positive probability, ξ_n avoids 0 for all $n \geq 1$.

Suppose not, so ξ a.s. revisits 0. Since particles in ξ reproduce independently, this implies that ξ returns to the origin infinitely often. This is a contradiction, as we showed the opposite in proving Proposition 18. \square

Proof of Theorem 5. In two steps, we bound the probability p of recurrence on $G = (\mathbb{Z} \cup \mathbb{T}_6)/\{0 \sim \emptyset\}$:

- ($p > 0$): The probability is 0 that any frog starting in \mathbb{Z} wakes but fails to visit 0, by the recurrence of simple random walk on \mathbb{Z} . All frogs at $1, 2, \dots$ wake up with positive probability by Lemma 21, and on this event they therefore all visit 0.
- ($p < 1$): With probability $6/8$ the first jump of the frog at $0 \in G$ will be into \mathbb{T}_6 . Conditional on this, Lemma 22 guarantees a frog model in this configuration will never again visit the origin with probability $p' > 0$. Therefore, $1 - p \geq \frac{6}{8}p'$. \square

5. CONJECTURES

Simulations suggest that for $d = 3$ the frog model on \mathbb{T}_d is recurrent a.s., while for $d = 4$ the model is transient a.s. Our approach was to consider the frog model with the addition of stunning fences at each depth. When a frog jumps on a fence for the first time, it is stunned and stops moving. When all frogs are stunned at depth k , the fence turns off, and frogs resume their motion until they reach depth $k + 1$ and are stunned again. Let $A_{d,k}$ be the number of stunned frogs that pile up on the fence at depth k before it turns off. We then examined the growth of $A_{d,k}$ in k for different choices of d . (The more obvious approach of directly simulating the frog model and counting visits to the root does not yield any obvious conclusions, as the rapid growth of the frog model makes it impossible to simulate very far.)

Martingale techniques tell us that the probability a frog at distance k from the root reaches the root before visiting depth $k + 1$ is greater than cd^{-k} for some $c > 0$ independent of k . It follows that

$$\mathbf{E}[\text{visits to root between } k\text{th and } (k+1)\text{th stunnings}] \geq cd^{-k} \mathbf{E}[A_{d,k}].$$

So, if $\sum_{k=1}^{\infty} d^{-k} \mathbf{E}[A_{d,k}] = \infty$, then the expected number of visits to the root is infinite, which strongly suggests the model is recurrent.

This occurs if $kd^{-k} \mathbf{E}[A_{d,k}]$ is bounded from below. The data in Figure 8 summarizes the behavior of $kd^{-k} \mathbf{E}[A_{d,k}]$ to the maximum k we could easily simulate, $k = 18$. For $d = 4$, the slow growth of $A_{d,k}$ makes us suspect that the model is transient. The plot for $d = 2$ confirms Theorem 1 (i). Interestingly, $d = 3$ appears to be recurrent but very near criticality. The different growth for $d = 3$ is grounds for further investigation: it is possible $d = 2$ and $d = 3$ exhibit different forms of recurrence.

For $d = 2$ the simulated values of $k2^{-k} \mathbf{E}[A_{2,k}]$ appear to be growing linearly. This suggests a constant expected number of returns between each successive stunning. As the average number of steps for an individual frog between stunnings is constant, this could indicate that the average time between returns is also bounded away from infinity. If this is the case then the probability that there is a frog at the origin at time t would be bounded away from 0 as t gets large. However, for $d = 3$ it appears that $k3^{-k} \mathbf{E}[A_{3,k}]$ is sublinear. This might indicate that the average time between returns is unbounded and the probability of a frog occupying the origin at time t is approaching 0 as t approaches infinity. This leads us to ask the following:




FIGURE 8. Plots of simulated values of $kd^{-k}A_{d,k}$ against k for $d = 2, 3, 4$. The number of simulations used in each estimate is shown in the chart.

Open Question 3. *Is the frog model on strongly recurrent on \mathbb{T}_2 but only weakly recurrent on \mathbb{T}_3 ?*

Such a result would have analogues with other interacting particle systems on trees. For example percolation on $\mathbb{T}_6 \times \mathbb{Z}$ has a three phases: no infinite components, infinitely many infinite components and a unique infinite component [GN90]. Similarly the contact process on trees can have strongly recurrent, weakly recurrent, and extinction phases [Pem92].

Acknowledgments. We would like to thank Shirshendu Ganguly for his suggestions throughout the project. Christopher Fowler helped with a calculation, Avi Levy asked a question which led to the inclusion of Lemma 11 and James Morrow helped address potential concerns about roundoff error. Thanks to Soumik Pal who pointed out for large d the dynamics should be simpler—this remark sparked our study of transience. We thank Robin Pemantle for directing our attention to [AB05]. We also thank Nina Gantert who pointed out an inaccuracy in a previous version.

The first author was partially supported by NSF grant DMS-1308645 and NSA grant H98230-13-1-0827, the second author by NSF CAREER award DMS-0847661 and NSF grant DMS-1401479, and the third author by NSF RTG grant 0838212.

REFERENCES

- [AB05] David J. Aldous and Antan Bandyopadhyay, *A survey of max-type recursive distributional equations*, Ann. Appl. Probab. **15** (2005), no. 2, 1047–1110. MR 2134098 (2007e:60010)
- [AMP02a] O. S. M. Alves, F. P. Machado, and S. Yu. Popov, *The shape theorem for the frog model*, Ann. Appl. Probab. **12** (2002), no. 2, 533–546. MR 1910638 (2003c:60159)
- [AMP02b] Oswaldo Alves, Fabio Machado, and Serguei Popov, *Phase transition for the frog model*, Electron. J. Probab. **7** (2002), no. 16, 1–21, <http://ejp.ejpecp.org/article/view/115>.
- [Big76] J. D. Biggins, *The first- and last-birth problems for a multitype age-dependent branching process*, Advances in Appl. Probability **8** (1976), no. 3, 446–459. MR 0420890 (54 #8901)
- [BR10] Jean Bérard and Alejandro F. Ramírez, *Large deviations of the front in a one-dimensional model of $X + Y \rightarrow 2X$* , Ann. Probab. **38** (2010), no. 3, 955–1018. MR 2674992 (2011e:60219)
- [BW03] Itai Benjamini and David B Wilson, *Excited random walk*, Electron. Comm. Probab **8** (2003), no. 9, 86–92.
- [CQR09] Francis Comets, Jeremy Quastel, and Alejandro F. Ramírez, *Fluctuations of the front in a one dimensional model of $X + Y \rightarrow 2X$* , Trans. Amer. Math. Soc. **361** (2009), no. 11, 6165–6189. MR 2529928 (2010i:60281)

- [DG99] D. J. Daley and J. Gani, *Epidemic modeling: an introduction*, Cambridge Studies in Mathematical Biology, vol. 15, Cambridge University Press, Cambridge, 1999. MR 1688203 (2000e:92042)
- [DP14] Christian Döbler and Lorenz Pfeifroth, *Recurrence for the frog model with drift on \mathbb{Z}^d* , Electron. Commun. Probab. **19** (2014), no. 79, 13. MR 3283610
- [DRS10] Ronald Dickman, Leonardo T. Rolla, and Vlasas Sidoravicius, *Activated random walkers: facts, conjectures and challenges*, J. Stat. Phys. **138** (2010), no. 1-3, 126–142. MR 2594894 (2011h:82047)
- [GN90] G. R. Grimmett and C. M. Newman, *Percolation in $\infty + 1$ dimensions*, Disorder in physical systems, Oxford Sci. Publ., Oxford Univ. Press, New York, 1990, pp. 167–190. MR 1064560 (92a:60207)
- [GNR15] Arka P. Ghosh, Steven Noren, and Alexander Roitershtein, *On the range of the transient frog model on \mathbb{Z}* , arXiv:1502.02738, 2015.
- [GS09] N. Gantert and P. Schmidt, *Recurrence for the frog model with drift on \mathbb{Z}* , Markov Process. Related Fields **15** (2009), no. 1, 51–58, <http://wwwmath.uni-muenster.de/statistik/gantert/frogs.pdf>. MR 2509423 (2010g:60170)
- [HJJ15] Chris Hoffman, Tobias Johnson, and Matthew Junge, *From transience to recurrence with Poisson tree frogs*, arXiv:1501.05874, 2015.
- [KPV04] Irina Kurkova, Serguei Popov, and M. Vachkovskaia, *On infection spreading and competition between independent random walks*, Electron. J. Probab. **9** (2004), no. 11, 293–315.
- [KS06] Harry Kesten and Vlasas Sidoravicius, *A phase transition in a model for the spread of an infection*, Illinois J. Math. **50** (2006), no. 1-4, 547–634. MR 2247840 (2007m:60298)
- [Liu98] Quansheng Liu, *Fixed points of a generalized smoothing transformation and applications to the branching random walk*, Adv. in Appl. Probab. **30** (1998), no. 1, 85–112. MR 1618888 (99f:60151)
- [LMP05] Élcio Levensztayn, Fábio Machado, and Serguei Popov, *An improved upper bound for the critical probability of the frog model on homogeneous trees*, Journal of Statistical Physics **119** (2005), no. 1-2, 331–345 (English), <http://dx.doi.org/10.1007/s10955-004-2051-8>.
- [Pem92] Robin Pemantle, *The contact process on trees*, Ann. Probab. **20** (1992), no. 4, 2089–2116. MR 1188054 (94d:60155)
- [Pem07] ———, *A survey of random processes with reinforcement*, Probab. Surv. **4** (2007), 1–79. MR 2282181 (2007k:60230)
- [Pop01] S.Yu. Popov, *Frogs in random environment*, Journal of Statistical Physics **102** (2001), no. 1-2, 191–201 (English).
- [Pop03] Serguei Yu. Popov, *Frogs and some other interacting random walks models*, Discrete random walks (Paris, 2003), Discrete Math. Theor. Comput. Sci. Proc., AC, Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2003, pp. 277–288 (electronic). MR 2042394
- [RS04] Alejandro F. Ramírez and Vlasas Sidoravicius, *Asymptotic behavior of a stochastic combustion growth process*, Journal of the European Mathematical Society **6** (2004), no. 3, 293–334.
- [RS12] Leonardo T. Rolla and Vlasas Sidoravicius, *Absorbing-state phase transition for driven-dissipative stochastic dynamics on \mathbb{Z}* , Invent. Math. **188** (2012), no. 1, 127–150. MR 2897694
- [SS07] Moshe Shaked and J. George Shanthikumar, *Stochastic orders*, Springer Series in Statistics, Springer, New York, 2007. MR 2265633 (2008g:60005)
- [ST14] Vlasas Sidoravicius and Augusto Teixeira, *Absorbing-state transition for Stochastic Sandpiles and Activated Random Walks*, arXiv:1412.7098, 2014.
- [TW99] András Telcs and Nicholas C. Wormald, *Branching and tree indexed random walks on fractals*, Journal of Applied Probability **36** (1999), no. 4, 999–1011.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WASHINGTON
 E-mail address: hoffman@math.washington.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF SOUTHERN CALIFORNIA
 E-mail address: tobias.johnson@usc.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WASHINGTON
 E-mail address: jungem@math.washington.edu